

When and How to Use FFT

The WP02 Spectral Analysis package, with FFT (Fast Fourier Transform), reveals signal characteristics not visible in the time domain and adds the power of frequency domain analysis to your oscilloscope. FFT converts a time domain waveform into frequency domain spectra similar to those of a spectrum analyzer, but with important differences and added benefits.

Why Use FFT?

For a large class of signals, greater insight can be gained by looking at spectral representation rather than time description. Signals encountered in the frequency response of amplifiers, oscillator phase noise and those in mechanical vibration analysis — to mention just some applications — are easier to observe in the frequency domain.

If sampling is done at a rate fast enough to faithfully approximate the original waveform (usually five times the highest frequency component in the signal), the resulting discrete data series will uniquely describe the analog signal.

This is of particular value when dealing with transient signals because, unlike FFT, conventional swept spectrum analyzers cannot handle them.

Theory Behind FFT

Spectral analysis theory assumes that the signal for transformation be of infinite duration. Since no physical signal can meet this condition, a useful assumption for reconciling theory and practice is to view the signal as consisting of an infinite series of replica of itself. These replica are multiplied by a rectangular window (the display grid) that is zero outside of the observation grid.

- *For an explanation of FFT terms: see the Glossary on page C-17*
- *Using FFT Functions: see page C-9*
- *FFT Algorithms: page C-14*

Figure 1 displays four panels related to signal analysis:

- Panel A:** Time-domain waveform showing a periodic signal. Scale: 2 ms, 26.5 mV.
- Panel B:** Power spectrum (PS) showing the frequency content. Scale: 1 kHz, 20.0 dBm, -13.994 dBm.
- Panel C:** FFT magnitude spectrum showing the frequency content. Scale: 2 ms, 26.5 mV AC, 1 V AC, 5 V AC, 5 V AC.
- Panel D:** Setup menu showing various parameters:
 - SETUP OF B
 - use Math? No **Yes**
 - Math Type: Enh. Res, Extrema, **FFT**, FFTAVG, Functions
 - FFT result: Phase, Power Dens, **Power Spect**, Real, Real+Imag
 - with window: Von Hann, DC, **AC**
 - of: **1**, 2, 3, 4, A, C, D
 - M1, M2, M3, M4
 - 2.5 MS/s
 - ☐ STOPPED

C-2

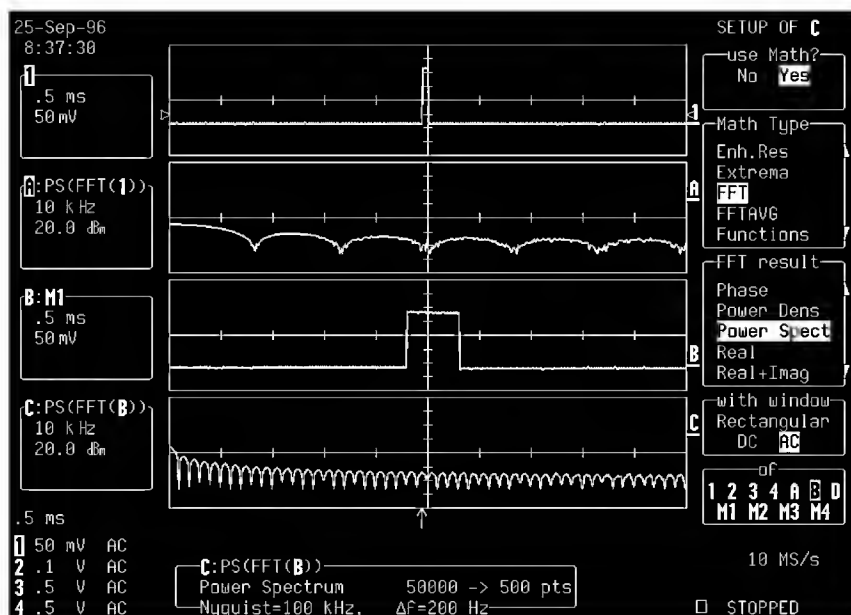


Figure C-2

An FFT operation on an N point time-domain signal may thus be compared to passing the signal through a comb filter consisting of a bank of $N/2$ filters. All the filters have the same shape and width and are centered at $N/2$ discrete frequencies. Each filter collects the signal energy that falls into the immediate neighborhood of its center frequency. Thus it can be said that there are $N/2$ frequency bins. The distance in Hz between the center frequencies of two neighboring bins is always the same: Δf .

Power (Density) Spectrum Because of the linear scale used to show magnitudes, lower amplitude components are often hidden by larger components. In addition to the functions offering magnitude and phase representations, the FFT option offers power density and power spectrum density functions, selected from the "FFT result" menu shown in the figures. These latter functions are even better suited for characterizing spectra. The power spectrum (V^2) is the square of the magnitude spectrum (0 dB m corresponds to voltage equivalent to 1 mW into 50 Ω .) This is the representation



of choice for signals containing isolated peaks — periodic signals, for instance.

The power density spectrum (V^2/Hz) is the power spectrum divided by the equivalent noise bandwidth of the filter in Hz associated with the FFT calculation. This is best employed for characterizing broad-band signals such as noise.

Memory for FFT

The amount of acquisition memory available will determine the maximum range (Nyquist frequency) over which signal components can be observed. Consider the problem of determining the length of the observation window and the size of the acquisition buffer if a Nyquist rate of 500 MHz and a resolution of 10 kHz are required. To obtain a resolution of 10 kHz, the acquisition time must be at least:

$$T = 1/\Delta f = 1/10 \text{ kHz} = 100 \mu\text{s}.$$

For a digital oscilloscope with a memory of 100 k, the highest frequency that can be analyzed is:

$$\Delta f \times N/2 = 10 \text{ kHz} \times 100 \text{ k}/2 = 500 \text{ MHz}.$$

FFT Pitfalls to Avoid

Take care to ensure that signals are correctly acquired: improper waveform positioning within the observation window produces a distorted spectrum. The most common distortions can be traced to insufficient sampling, edge discontinuities, windowing or the “picket fence” effect.

Because the FFT acts like a bank of bandpass filters centered at multiples of the frequency resolution, components that are not exact multiples of that frequency will fall within two consecutive filters. This results in an attenuation of the true amplitude of these components.

Picket Fence and Scallop

The highest point in the spectrum can be 3.92 dB lower when the source frequency is halfway between two discrete frequencies. This variation in spectrum magnitude is the picket fence effect. And the corresponding attenuation loss is referred to as scallop loss. LeCroy scopes automatically correct for the scallop effect, ensuring that the magnitude of the spectra lines correspond to their true values in the time domain.

If a signal contains a frequency component above Nyquist, the spectrum will be aliased, meaning that the frequencies will be folded back and spurious. Spotting aliased frequencies is often difficult, as the aliases may ride on top of real harmonics. A simple way of checking is to modify the sample rate and verify whether the frequency distribution changes.

Leakage

FFT assumes that the signal contained within the time grid is replicated endlessly outside the observation window. Therefore if the signal contains discontinuities at its edges, pseudo-frequencies will appear in the spectral domain, distorting the real spectrum. When the start and end phase of the signal differ, the signal frequency falls within two frequency cells broadening the spectrum.

This effect is illustrated in Figure C-1. Because the display does not contain an integral number of periods, the spectrum displayed in Trace B does not reveal sharp frequency components. Intermediate components exhibit a lower and broader peak. The broadening of the base, stretching out in many neighboring bins, is termed leakage. Cures for this are to ensure that an integral number of periods is contained within the display grid or that no discontinuities appear at the edges. Another is to use a window function to smooth the edges of the signal.

Choosing a Window

The choice of a spectral window is dictated by the signal's characteristics. Weighting functions control the filter response shape and affect noise bandwidth as well as side-lobe levels. Ideally, the main lobe should be as narrow and flat as possible to effectively discriminate all spectral components, while all side lobes should be infinitely attenuated.

Chosen from the "with window" menu, the window type defines the bandwidth and shape of the equivalent filter to be used in the FFT processing.

In the same way as one would choose a particular camera lens for taking a picture, some experimenting is generally necessary to determine which window is most suitable. However, the following general guidelines should help (*see page C-11 for window types*).



Rectangular windows provide the highest frequency resolution and are thus useful for estimating the type of harmonics present in the signal. Because the rectangular window decays as a $\sin x/x$ function in the spectral domain, slight attenuation will be induced. Alternative functions with less attenuation — Flat-top and Blackman-Harris — provide maximum amplitude at the expense of frequency resolution. Whereas, Hamming and von Hann are good for general purpose use with continuous waveforms.

Improving Dynamic Range Enhanced resolution (*see Appendix B*) uses a low pass filtering technique that can potentially provide for three additional bits (18 dBs) if the signal noise is uniformly distributed (white). Low pass filtering should be considered when high frequency components are irrelevant. A distinct advantage of this technique is that it works for both repetitive and transient signals. The SNR increase is conditioned by the cut-off frequency of the Eres low pass filter and the noise shape (frequency distribution).

LeCroy digital oscilloscopes employ FIR digital filters so that a constant phase shift is maintained. The phase information is therefore not distorted by the filtering action.

Spectral Power Averaging Even greater dynamic-range improvement is obtained on signals showing periodicity. Moreover, the range can be increased without sacrificing frequency response. The LeCroy oscilloscope being used is equipped with accumulation buffers 32 bits wide to prevent overflows.

Spectral power averaging is useful when the signal varies in time and the mean power of the signal needs to be estimated. Typical applications include noise and pseudo- random noise. Whereas time averaging ignores phase information, spectral averaging tracks magnitude as well as phase information. It is thus a superior estimator. And the improvement is typically proportional to the square root of the number of averages. For instance, averaging white noise at full scale over 10 sweeps yields a typical improvement of nearly 20 dBs.

Spectral power averaging is the technique of choice when determining the frequency response of passive networks such as filters. Figures 3 and 4 show the transfer functions of a low pass filter with a 3 dB cutoff of 11 MHz obtained by exciting the filter with a white noise source (Fig. C-3) and a sine swept generator (Fig. C-4). Both techniques give substantially the same results. The choice of method is governed by the availability of an adequate generating source.

The spectra of single time-domain waveforms can be computed and displayed to obtain power averages obtained over as many as 50 000 spectra.




Figure C-3




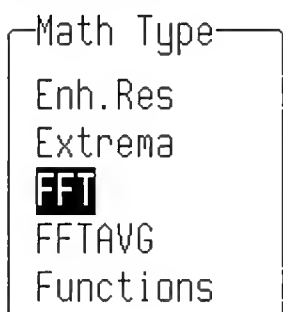
Figure C-4

Overall...

Because of its versatility, FFT analysis has become a popular analysis tool. However, some care must be taken with it. In most instances, incorrect positioning of the signal within the display grid will significantly alter the spectrum. Effects such as leakage and aliasing that distort the spectrum must be understood if meaningful conclusions are to be arrived at when using FFT.

An effective way to reduce these effects is to maximize the acquisition record length. Record length directly conditions the effective sampling rate of the scope and therefore determines the frequency resolution and span at which spectral analysis can be carried out.

Using FFT Functions



Select "FFT" from the "Math Type" menu (see Chapter 10 for a full description of math and waveform processing menus). Spectra displayed with a linear frequency axis running from zero to the Nyquist frequency are shown at the right-hand edge of the trace. The frequency scale factors (Hz/div) are in a 1–2–5 sequence.

The processing equation is displayed at the bottom of the screen, together with the three key parameters that characterize an FFT spectrum. These are:

1. Transform Size N (number of input points)
2. Nyquist frequency (= $\frac{1}{2}$ sample rate), and
3. Frequency Increment, Δf , between two successive points of the spectrum.

These parameters are related as:

$$\text{Nyquist frequency} = \Delta f * N/2.$$

Where: $\Delta f = 1/T$, and where T is the duration of the input waveform record ($10 * \text{time/div}$). The number of output points is equal to $N/2$.

Note on Maximum Points: FFT spectra are computed over the entire source time-domain waveform. This limits the number of points used for FFT processing. If the input waveform contains more points than the selected maximum (in "for Math use max points", they are decimated before FFT processing. But if it has fewer, all points are used.



The following selections can be made using the "FFT result" menu.

Phase

Measured with respect to a cosine whose maximum occurs at the left-hand edge of the screen, at which point it has 0°. Similarly, a positive-going sine starting at the left-hand edge of the screen has a -90° phase. (Displayed in degrees.)

Power Density

The signal power normalized to the bandwidth of the equivalent filter associated with the FFT calculation. The power density is suitable for characterizing broad-band noise. (It is displayed on a logarithmic vertical axis calibrated in dBm.)

Power Spectrum

The signal power (or magnitude) represented on a logarithmic vertical scale: 0 dBm corresponds to the voltage (0.316 V peak) which is equivalent to 1 mW into 50 Ω. The power spectrum is suitable for characterizing spectra which contain isolated peaks. (dBm.)

Magnitude

The peak signal amplitude represented on a linear scale. (Same units as input signal.)

Real, Real + Imaginary, Imaginary

These represent the complex result of the FFT processing. (Same units as input signal.)


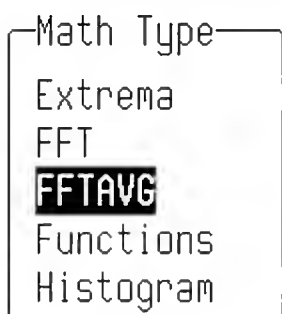
Windows

Chosen using the “with window” menu, the window type defines the bandwidth and shape of the filter to be used in the FFT processing (*see the table on page C-17 for these filters’ parameters*). When “AC” is selected from the same menu, the DC component of the input signal is forced to zero prior to the FFT processing. This improves the amplitude resolution, especially when the input has a large DC component.

Window Type	Applications and Limitations
Rectangular	Normally used when the signal is transient — completely contained in the time-domain window — or known to have a fundamental frequency component that is an integer multiple of the fundamental frequency of the window. Signals other than these types will show varying amounts of spectral leakage and scallop loss, corrected by selecting another type of window.
Hanning (Von Hann)	Reduce leakage and improve amplitude accuracy. However, frequency resolution is also reduced.
Hamming	Reduce leakage and improve amplitude accuracy. However, frequency resolution is also reduced.
Flat Top	This window provides excellent amplitude accuracy with moderate reduction of leakage, but also at the loss of frequency resolution.
Blackman–Harris	It reduces the leakage to a minimum, but again along with reduced frequency resolution.

FFT Power Average

A function can be defined as the power average of FFT spectra computed by another function (*see page C-6*). Choose “FFTAVG” from the “Math Type” Menu, and “Power Spect” from “FFT Result”.





Additional Processing

Other waveform processing functions, such as Averaging and Arithmetic, can be applied to waveforms before FFT processing is performed. Time-domain averaging prior to FFT, for example, can be used if a stable trigger is available to reduce random noise in the signal.

Note:

- *To increase the FFT frequency range, the Nyquist frequency, raise the effective sampling frequency by increasing the maximum number of points or using a faster time base.*
- *To increase the FFT frequency resolution, increase the length of the time-domain waveform record by using a slower time base.*

Memory Status

When FFT is used, the field beneath the grid displays parameters of the waveform descriptor, including number of points, horizontal and vertical scale factors and units.

C:PS(AVGP(B))

Power Spectrum 10000 -> 2500 pts

Using Cursors with FFT

For reading the amplitude and frequency of a data point, the Absolute Time cursor can be moved into the frequency domain by going beyond the right-hand edge of a time-domain waveform.

The Relative Time cursors can be moved into the frequency domain to simultaneously indicate the frequency difference and the amplitude difference between two points on each frequency-domain trace.

The Absolute Voltage cursor reads the absolute value of a point in a spectrum in the appropriate units, and the *Relative Voltage* cursors indicate the difference between two levels on each trace

Error Messages

One of these FFT-related error messages may be displayed at the top of the screen.

Error Messages	
Message	Meaning
"Incompatible input record type"	FFT power average is defined only on a function defined as FFT.
"Horizontal units don't match"	FFT of a frequency-domain waveform is not available.
"FFT source data zero filled"	If there are invalid data points in the source waveform (at the beginning or at the end of the record), these are replaced by zeros before FFT processing.
"FFT source data over/underflow"	The source waveform data has been clipped in amplitude, either in the acquisition — gain too high or inappropriate offset — or in previous processing. The resulting FFT contains harmonic components which would not be present in the unclipped waveform. The settings defining the acquisition or processing should be changed to eliminate the over/underflow condition.
"Circular computation"	A function definition is circular (i.e. the function is its own source, indirectly via another function or expansion). One of the definitions should be changed.



FFT Algorithms

A summary of the algorithms used in the oscilloscope's FFT computation is given here in the form of seven steps:

1. If the maximum number of points is smaller than the source number of points, the source waveform data are decimated prior to the FFT. These decimated data extend over the full length of the source waveform. The resulting sampling interval and the actual transform size selected provide the frequency scale factor in a 1–2–5 sequence.
2. The data are multiplied by the selected window function.
3. FFT is computed, using a fast implementation of the DFT (Discrete Fourier Transform):

$$X_n = \frac{1}{N} \sum_{k=0}^{N-1} x_k \times W^{nk},$$

where: x_k is a complex array whose real part is the modified source time domain waveform, and whose imaginary part is 0; X_n is the resulting complex frequency-domain waveform; $W = e^{-2\pi j/N}$; and N is the number of points in x_k and X_n .

The generalized FFT algorithm, as implemented here, works on N , which need *not* be a power of 2.

4. The resulting complex vector X_n is divided by the coherent gain of the window function, in order to compensate for the loss of the signal energy due to windowing. This compensation provides accurate amplitude values for isolated spectrum peaks.
5. The real part of X_n is symmetric around the Nyquist frequency, i.e.

$$R_n = R_{N-n},$$

while the imaginary part is asymmetric, i.e.

$$I_n = -I_{N-n}.$$

The energy of the signal at a frequency n is distributed equally between the first and the second halves of the spectrum; the energy at frequency 0 is completely contained in the 0 term.

The first half of the spectrum (Re, Im), from 0 to the Nyquist frequency is kept for further processing and doubled in amplitude:

$$R'_n = 2 \cdot R_n \quad 0 \leq n < N/2$$

$$I'_n = 2 \cdot I_n \quad 0 \leq n < N/2.$$

6. The resultant waveform is computed for the spectrum type selected.

If “Real”, “Imaginary”, or “Real + Imaginary” is selected, no further computation is needed. The appropriate part of the complex result is given as the result (R'_n or I'_n or $R'_n + jI'_n$, as defined above).

If “Magnitude” is selected, the magnitude of the complex vector is computed as:

$$M_n = \sqrt{R_n'^2 + I_n'^2}.$$

Steps 1–6 lead to the following result:

An AC sine wave of amplitude 1.0 V with an integral number of periods N_p in the time window, transformed with the rectangular window, results in a fundamental peak of 1.0 V magnitude in the spectrum at frequency $N_p \times \Delta f$. However, a DC component of 1.0 V, transformed with the rectangular window, results in a peak of 2.0 V magnitude at 0 Hz.

The waveforms for the other available spectrum types are computed as follows:

$$\begin{aligned} \text{Phase: angle} &= \arctan(I_n/R_n) & M_n > M_{min} \\ \text{angle} &= 0 & M_n \leq M_{min}. \end{aligned}$$

Where M_{min} is the minimum magnitude, fixed at about 0.001 of the full scale at any gain setting, below which the angle is not well defined.

The dBm Power Spectrum:

$$dBm PS = 10 \times \log_{10} \left(\frac{M_n^2}{M_{ref}^2} \right) = 20 \times \log_{10} \left(\frac{M_n}{M_{ref}} \right)$$

where $M_{ref} = 0.316$ V (that is, 0 dBm is defined as a sine wave of 0.316 V peak or 0.224 V RMS, giving 1.0 mW into 50Ω).

The dBm Power Spectrum is the same as dBm Magnitude, as suggested in the above formula.

dBm Power Density:

$$dBm PD = dBm PS - 10 \times \log_{10} (ENBW \times \Delta f)$$

where $ENBW$ is the equivalent noise bandwidth of the filter corresponding to the selected window, and Δf is the current frequency resolution (bin width).

7. The FFT Power Average takes the complex frequency-domain data R'_n and I'_n for each spectrum generated in Step 5, and computes the square of the magnitude:

$$M_n^2 = R_n'^2 + I_n'^2,$$

then sums M_n^2 and counts the accumulated spectra. The total is normalized by the number of spectra and converted to the selected result type using the same formulae as are used for the Fourier Transform.

Glossary

Defines the terms frequently used in FFT spectrum analysis and relates them to the oscilloscope.

Aliasing

If the input signal to a sampling acquisition system contains components whose frequency is greater than the Nyquist frequency (half the sampling frequency), there will be less than two samples per signal period. The result is that the contribution of these components to the sampled waveform is indistinguishable from that of components below the Nyquist frequency. This is **aliasing**.

The timebase and transform-size should be selected so that the resulting Nyquist frequency is higher than the highest significant component in the time-domain record.

Coherent Gain

The normalized coherent gain of a filter corresponding to each window function is 1.0 (0 dB) for a rectangular window and less than 1.0 for other windows. It defines the loss of signal energy due to the multiplication by the window function. This loss is compensated in the oscilloscope. This table lists the values for the implemented windows.

Window Frequency-Domain Parameters				
Window Type	Highest Side Lobe (dB)	Scallop Loss (dB)	ENBW (bins)	Coherent Gain (dB)
Rectangular	-13	3.92	1.0	0.0
von Hann	-32	1.42	1.5	-6.02
Hamming	-43	1.78	1.37	-5.35
Flat Top	-44	0.01	2.96	-11.05
Blackman-Harris	-67	1.13	1.71	-7.53



ENBW

Equivalent Noise BandWidth (ENBW) is the bandwidth of a rectangular filter (same gain at the center frequency), equivalent to a filter associated with each frequency bin, which would collect the same power from a white noise signal. In the table on the previous page, the ENBW is listed for each window function implemented and is given in bins.

Filters

Computing an N-point FFT is equivalent to passing the time-domain input signal through N/2 filters and plotting their outputs against the frequency. The spacing of filters is $\Delta f = 1/T$ while the bandwidth depends on the window function used (see Frequency bins).

Frequency bins

The FFT algorithm takes a discrete source waveform, defined over N points, and computes N complex Fourier coefficients, which are interpreted as harmonic components of the input signal.

For a real source waveform (imaginary part equals 0), there are only N/2 independent harmonic components.

An FFT corresponds to analyzing the input signal with a bank of N/2 filters, all having the same shape and width, and centered at N/2 discrete frequencies. Each filter collects the signal energy that falls into the immediate neighborhood of its center frequency, and thus it can be said that there are N/2 “frequency bins”.

The distance in hertz between the center frequencies of two neighboring bins is always:

$$\Delta f = 1/T,$$

where T is the duration of the time-domain record in seconds.

The width of the main lobe of the filter centered at each bin depends on the window function used. The rectangular window has a nominal width at 1.0 bin. Other windows have wider main lobes (*see table*).

Frequency Range

The range of frequencies computed and displayed is 0 Hz (displayed at the left-hand edge of the screen) to the Nyquist frequency (at the rightmost edge of the trace).

Frequency Resolution

In a simple sense, the frequency resolution is equal to the bin width Δf . That is, if the input signal changes its frequency by Δf , the corresponding spectrum peak will be displaced by Δf . For smaller changes of frequency, only the shape of the peak will change.

However, the effective frequency resolution (i.e. the ability to resolve two signals whose frequencies are almost the same) is further limited by the use of window functions. The ENBW value of all windows other than the rectangular is greater than Δf and the bin width. *The table on page C-17 lists the ENBW values for the implemented windows.*

Leakage

In the power spectrum of a sine wave with an integral number of periods in the (rectangular) time window (i.e. the source frequency equals one of the bin frequencies), the spectrum contains a sharp component whose value accurately reflects the source waveform's amplitude. For intermediate input frequencies this spectral component has a lower and broader peak.

The broadening of the base of the peak, stretching out into many neighboring bins is termed *leakage*. It is due to the relatively high side lobes of the filter associated with each frequency bin.

The filter side lobes and the resulting leakage are reduced when one of the available window functions is applied. The best reduction is provided by the Blackman–Harris and Flat Top windows. However, this reduction is offset by a broadening of the main lobe of the filter.

Number of Points

FFT is computed over the number of points (Transform Size) whose upper bounds are the source number of points, and by the maximum number of points selected in the menu. FFT generates spectra of $N/2$ output points.

Nyquist Frequency

The Nyquist frequency is equal to one half of the effective sampling frequency (after the decimation): $\Delta f \times N/2$.

Picket Fence Effect

If a sine wave has a whole number of periods in the time domain record, the power spectrum obtained with a rectangular window will have a sharp peak, corresponding exactly to the frequency and amplitude of the sine wave. Otherwise the spectrum peak with a rectangular window will be lower and broader.

The highest point in the power spectrum can be 3.92 dB lower (1.57 times) when the source frequency is halfway between two



discrete bin frequencies. This variation of the spectrum magnitude is called the *picket fence effect* (the loss is called the scallop loss).

All window functions compensate this loss to some extent, but the best compensation is obtained with the Flat Top window.

Power Spectrum

The power spectrum (V^2) is the square of the magnitude spectrum.

The power spectrum is displayed on the dBm scale, with 0 dBm corresponding to:

$$V_{\text{ref}}^2 = (0.316 V_{\text{peak}})^2,$$

where V_{ref} is the peak value of the sinusoidal voltage, which is equivalent to 1 mW into 50 Ω .

Power Density Spectrum

The power density spectrum (V^2/Hz) is the power spectrum divided by the equivalent noise bandwidth of the filter in hertz. The power density spectrum is displayed on the dBm scale, with 0 dBm corresponding to ($V_{\text{ref}}^2/\text{Hz}$).

Sampling Frequency

The time-domain records are acquired at sampling frequencies dependent on the selected time base. Before the FFT computation, the time-domain record may be decimated. If the selected maximum number of points is lower than the source number of points, the effective sampling frequency is reduced. The effective sampling frequency equals twice the Nyquist frequency.

Scallop Loss

Loss associated with the picket fence effect.

Window Functions

All available window functions belong to the sum of cosines family with one to three non-zero cosine terms:

$$W_k = \sum_{m=0}^{M-1} a_m \cos \left(\frac{2\pi k}{N} m \right) \quad 0 \leq k < N,$$

where: $M=3$ is the maximum number of terms, a_m are the coefficients of the terms, N is the number of points of the decimated source waveform, and k is the time index.

The following table lists the coefficients a_m . The window functions seen in the time domain are symmetric around the point $k = N/2$.

Coefficients Of Window Functions			
Window Type	a0	a1	a2
Rectangular	1.0	0.0	0.0
von Hann	0.5	-0.5	0.0
Hamming	0.54	-0.46	0.0
Flat-Top	0.281	-0.521	0.198
Blackman-Harris	0.423	-0.497	0.079

Appendix C References

Bergland, G.D., *A Guided Tour of the Fast Fourier Transform*, IEEE Spectrum, July 1969, pp. 41–52.

A general introduction to FFT theory and applications.

Brigham, E.O., *The Fast Fourier Transform*, Prentice Hall, Inc., Englewood Cliffs, N.J., 1974.

Theory, applications and implementation of FFT. Includes discussion of FFT algorithms for N not a power of 2.

Harris, F.J., *On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform*, Proceedings of the IEEE, vol. 66, No. 1, January 1978, pp. 51–83.

Classic paper on window functions and their figures of merit, with many examples of windows.

Ramirez, R.W., *The FFT Fundamentals and Concepts*, Prentice Hall, Inc., Englewood Cliffs, N.J., 1985.

Practice-oriented, many examples of applications.